EUI WORKING PAPERS
ECO No. 2006/24

Fair Divisions as Attracting Nash Equilibria of Simple Games

MARCO GALBIATI

EUROPEAN UNIVERSITY INSTITUTE
Department of Economics

EUROPEAN UNIVERSITY INSTITUTE
DEPARTMENT OF ECONOMICS

Fair Divisions as Attracting Nash Equilibria of Simple Games

MARCO GALBIATI

EUI Working Paper **ECO** No. 2006/24

ISSN 1725-6704

Printed in Italy
European University Institute
Badia Fiesolana
I – 50016 San Domenico di Fiesole (FI)
Italy

http://www.iue.it/
http://cadmus.iue.it/dspace/index.jsp

Fair Divisions as Attracting Nash Equilibria of Simple Games

Marco Galbiati, IUE Florence*

Abstract

We consider the problem of allocating a finite number of divisible homogeneous goods to $N \geq 2$ individuals, in a way which is both envy-free and Pareto optimal. Building on Thomson (2005 Games and Economic Behavior), a new simple mechanism is presented here with the following properties: a) the mechanism fully implements the desired divisions, i.e. for each preference profile the set of equilibrium outcomes coincides with the set of fair divisions; b) the set of equilibria is a global attractor for the best-reply dynamics. Thus, players myopically adapting their strategies settle down in an fair division. The result holds even if mixed strategies are used.

Keywords:
Fair divisions, envy-free, implementation, best reply dynamics.
JEL Codes: C78, C73.

*This paper is part of my PhD dissertation: I thank K. Schag for his suggestions. I thank W. Thomson for providing me with a (then) unpublished version of his "Divide and Permute". I am also grateful to S. Brams and L. Samuelson for valuable suggestions. Discussions with A. Villar guided me through the literature at an early stage of this paper. Precious feedback came from participants to conferences where the paper was presented: the Second Game Theory Society Congress 2004, Marseille; the Econometric Society European Meeting, 2004 Madrid, Theory Lunches at the University of Madison, WI. All errors are mine. Author contact: marco.galbiati@iue.it.

1 Introduction

Consider a set of individuals who have to divide a bundle of homogeneous divisible goods among themselves. A referee, ignorant of the preferences of the individuals, wants the players to end up in a fair division, defined as a division which is both envy-free (EF) and Pareto optimal (PO). Envy-freeness requires that nobody prefers somebody else's share to her own. Pareto optimality excludes divisions like "everybody gets nothing", which are trivially envy-free but certainly unappealing. Although alternative definitions of fairness are possible, envy-freeness and efficiency have played a prominent role in the fair division literature. The appeal of these properties is probably due to the fact that, beside being intuitive, they refer to ordinal preferences and they do not involve interpersonal utility comparisons.

Our fair division problem can be approached in different ways; to begin with, it can be seen as a theoretical implementation problem. One can prove that under the usual assumptions on preferences (strict monotonicity, continuity, convexity), the fair division correspondence is non-empty and satisfies monotonicity and (vacously) no-veto power. As a consequence, our correspondence is for example implemented by Maskin's (1977, 2002) classic game form, when there are at least 3 players. As widely recognized however, this literature aims at probing the theoretical bounds of implementation. Indeed, most mechanisms in this pure implementation literature are aimed at proving the implementability of classes of rules, rather than at providing a workable solution to applied problems. As a consequence these mechanisms usually suffer from two practical problems: first, it is unlikely that real persons can manage their large strategy spaces. Second, most of the existing mechanisms are essentially coordination games with multiple equilibria; it is not clear how players can coordinate on any one of them. We return on this point later on.

Beside the implementation literature there exists another line of research, which follows a "procedural approach" to the fair division problem. Instead of formally defining game forms and adopting precise equilibrium concepts, these contributions give "protocols", mean to lead real claimants to a division. Results along this line range from evolutions of a divide-and-choose procedure by Banach and Knaster, to methods of sequential allocation. For an authoritative overview see Brams and Taylor (1996). These mechanisms are meant to be as simple as possible in order to serve concrete cases of division. Most of them *could* be formalized as game forms, but their equilibria are not often investigated. Instead, players are assumed to use safety strategies. There are of course some reasons for this: a) in many cases, manipulation of the mechanisms would be difficult, b) the Nash equilibria are so difficult to calculate that they don't seem a good predictor of play, c) the important properties of safety (minmax) strategies make them an appealing solution concept. To sum up, on one hand this literature yields mechanisms that are reasonably manipulation-proof, workable and elegant. On the other, it lacks a full game-theoretic formalization. More importantly, there is no known procedure in this literature to yield divisions which are at the same time PO and EF.[1]

Ideally located between the implementation literature and the procedural approach, there is the "Divide and Permute" mechanism (Thomson (2005)). This is a formally defined game form, which implements fair divisions in pure Nash equilibria (henceforth pNE), and which is also simple enough to be applied to real cases. However, Divide and Permute suffers from some important drawbacks, which indeed motivate our paper. Before illustrating these limitations, we briefly comment on "simplicity" in mechanism design, which is Divide and Permute's most appealing quality.

The issue of simplicity is in mechanism design is controversial because simplicity itself is an elusive concept. However, the existing literature seems to suggests that Divide and Permute (and so our games) are, in some sense, the simplest possible ones for the problem in hand. Dutta et al (1995) make it clear that, to implement Pareto efficient allocations, in equilibrium the mechanism must reveal the marginal rates of substitutions among goods. Thus, we cannot hope to solve our problem without players announcing prices, or something *de facto* equivalent. On the other hand, Saijo et al.(1996) suggest four properties as a definition of *natural* (i.e. simple) mechanisms: finite dimension (of strategy spaces), feasibility (for any strategy profile, outcomes respect a budget balance), best response (each player has a best reply against any strat-

[1] A famous mechanism, due to Brams and Taylor (1995), is for EF only.

egy profile) and forthrightedness (in equilibrium each players receives what she announced for herself). The canonic Maskin's mechanism for example, violates the first conditions, requiring players to announce preference profiles. On the contrary, our mechanisms possess the first three properties and essentially the forth one. Apart from Divide and Permute, there is one other "simple" game to implement fair allocations, by Saijo et al.(1999). We comment later on this mechanism; now we illustrate a major shortcoming of this game, limitation also present in Divide and Permute.

Thomson's (1999) and Saijo et al.'s (1999) game forms are a good solution to the problem of complexity, but they suffer from the other drawback common to most implementation literature: being coordination games, these mechanisms present an equilibrium selection problem. This feature is bound to emerge whenever one sets about to implement a generically multi-valued rule. What happens is that, for example, there is one equilibrium for each fair division. These latter are many and cannot be ranked in the same way by all players, so it is not clear which one should emerge, if ever. Althugh inavoidable, this problem must be dealt with. Achieving fair divisions is about enabling people to agree, so if we posit that our players are able to coordinate, we assume the problem away rather than solving it. Moreover, the coordination problem is worsened when the outcome function is discontinuous. As already pointed out by Postlewaite and Wettstein (1989), discontinuity is highly regrettable in implementation, because little mistakes on the part of players may imply the target to be missed by much. It is not clear if the continuity problem can be solved in game forms implementing EF and PO allocations. Up to now, there are no mechanisms with this property.

In this paper, the coordination problem is overcome by a learning argument. In addition to full implementation, we require that the equilibria of a mechanism be limit points of a dynamic adjustment process. More precisely, the game form presented here has the following features a) its strategy spaces and outcome functions are simple in the sense of Saijo et al. (1996); b) it fully implements the fair division rule in *pure* strategy Nash equilibria; c) for a version of perturbed best reply similar to that in Cabrales (1999), its outcome converge with probability one to an ε-equilibrium. In turn, the set of

ε-equilibria is a neighborood of the set of fair divisions; thus the limit outcomes of our game are ε-fair, in a sense to be made precise later on. In a word: we have *implementation* of fairness in Nash equilibrium, and *dynamic implementation* of ε-fairness in ε-equilibrium. It will finally be noted that those mixed-strategy equilibria which produce non-fair divisions are weeded out by the dynamics.

Recent works (Cabrales 1999, Cabrales and Ponti 2000) find some convergence results in classic mechanisms.[2] In particular Cabrales (1999) shows that discrete best reply dynamics, applied to Maskin's (1977) game form, singles out stable Nash equilibria. However, as we argued above, these general game forms are not suitable for concrete fair division problems. Being concerned with the issue of simplicity, our game seems a better solution to the problem in hand.

The next Section 2 lays down some notation and states the division problem in formal terms. Following Thomson (2005), from which this work evidently draws, Section 3 presents three game forms. The first two give respectively EF and PO divisions. The third one, combination of the previous two, implements divisions which are at the same time EF and PO. Section 4 deals with the dynamic properties of this final game form and Section 5 concludes.

2 Notation and general setting

We have an endowment of l homogeneous divisible goods, to be divided among N individuals. No restriction is imposed on N but finiteness. The endowment is represented by a vector $\omega \in \mathbb{R}^l_+$. Players have strictly monotonic, continuous, convex preferences over own bundles of goods, i. e. over vectors $z \in \mathbb{R}^l_+$. Individual preferences are occasionally described with the symbols $\succeq_i$, with $\succeq$ representing a profile of preferences. The set of all $\succeq$s satisfying strict monotonicity, continuity, convexity is indicated with R.

[2] Cabrales (1999) analyzes a slightly modified version of Maskin's (1977) game form as presented by Repullo (1987). Cabrales and Ponti (2000) consider Sjöström (1994) mechanism.

We define Z as the set of partitions of ω among the N players:

$$Z = \left\{ z = (z_1 ... z_N) \in \left(\mathbb{R}^l_+ \right)^N : \Sigma z_i = \omega \right\}$$

Given a preference profile $\succeq \in R$, the set of efficient divisions is defined as:

$$PO(\succeq) = \{ z \in Z : z \text{ is Pareto optimal in } Z \text{ for } \succeq \}$$

Instead, envy-free divisions are allocations whereby no player prefers somebody else's share to his own:

$$EF(\succeq) = \{ z \in Z : z_i \succeq_i z_j \forall i, j \}$$

Finally, the set of fair divisions is

$$F(\succeq) = EF(\succeq) \cap PO(\succeq)$$

The correspondence $F : R \to Z$ is the "fair division rule". It is well known that$F(\succeq)$ is non-empty for any $\succeq$ in R.[3] It is also easy to show by means of continuity arguments, that F is generically a proper (non-single valued) correspondence.

We now cast the fair division problem in the framework of implementation theory. A game form for our problem is a couple $\Gamma = \langle S, h \rangle$ such that: $S = \times S_i$ is some product strategy space, and g is an outcome function $g : S \to Z$. Preferences over bundles naturally define preferences over partitions, for which we use the symbol $\succeq$ again to simplify notation: for $x, y \in Z$ we have $x \succeq_i y$ when $x_i \succeq_i y_i$. Preferences over bundles thus define preferences over divisions and hence over outcomes of Γ. Given a preference profile, game form Γ then becomes a properly defined game $\Gamma' = \langle S, h, \succeq \rangle$. Consider now the set of pure Nash equilibria of Γ', which we indicate with $pNE(\succeq)$ (or simply pNE), and the corresponding set of equilibrium outcomes $g[pNE(\succeq)]$. Game form Γ fully implements F in pure Nash equilibria when $g[pNE(\succeq)] = F(\succeq)$ for every $\succeq \in R$.

One caveat: the sets EF, PO, F are dependent on the preference profile, and so is the set pNE for a given a mechanism. For correctness, explicit

[3] Perhaps Varian (1974) noticed this first. For example, a competitive equilibrium where all player's bundles have the same value is a fair allocation.

reference to the preference profile should be maintained. However, all proofs in this paper will *evidently* go through whatever the chosen $\succeq \in R$. This is different from saying that the statements are themselves evident, but it spares us pedantic references to $\succeq$ every time we talk about equilibria. Thus for example, *when we have the term "pNE" in a statement, this actually stands for "$pNE(\succeq)$" and the statement will be valid for all* $\succeq \in R$ -which is the kind of result we need for implementation.

Our problem is twofold. First of all, we want to fully implement the correspondence F in Nash equilibria by some mechanism Γ. Then, we want the set of equilibria to be an attractor for some dynamics, to be described later on.

3 Viable fair division games

This section presents three game forms implementing respectively envy-free, efficient and fair divisions. Our mechanisms evidently draw on Thomson (2005); they are an improvement over this latter because: i) they feature symmetric strategy spaces ii) their equilibria are learned by myopic players, as shown in Section 4. Following most of the implementation literature (see e.g. Maskin et al 2002 and references therein), we consider pure strategies only. This restriction is innocuous, as discussed in Section 4.

3.1 Envy freeness: Γ^{EF}

The first game implements EF divisions. Like all other games in this paper, it is a one-shot game. We first describe it informally as if it were a sequential game, to clarify its logic.

Each of the N players suggests an allocation and one of these proposals, say z, is selected (we see in a moment how). Then, the mechanism enables *each* player to choose his favourite share in z, whatever others' strategies. Thus, in equilibrium it will be the case that *all* players receive their favourite shares in z. More in detail, how is z selected? Beside suggesting an allocation, each player i names an integer and another player $k^i \neq i$. The integers are fed fed into a modulo game, and thus select one player (the "winner", w). Then, the "reference division" z is the one suggested by the player indicated by winner: $z = z^{k^w}$. On the other hand, how can players choose any share in z^{k^w}? Every i announces a permutation of the shares (π^i); these permutations are then applied to z^{k^w} in any predetermined order, before the resulting allocation is finally given out. It is clear that by playing an opportune π^i, every i can to reshuffle the shares in his favour, obtaining his preferred share in z^{k^w}.

Note that the winner's choice of z is limited by the fact that $k^w \neq w$. Similarly, we must also restrain the choice of k^w in selecting the reference allocation z.[4] We do so imposing a punishment on the loosers (so on k^w as well) in case $z^{k^w} \neq z^w$. By so doing, it will turn out that in equilibrium each player can pick any share in any of the proposals $\{z^i\}^{i=1..N}$, which in turn ensures envy-freeness (actually even more).

With a slight abuse of notation, let us write $N = \{1, 2, ..N\}$ and define Π the set of permutations $N \to N$. A formal description of the game follows.

Game form Γ^{PO} is $\langle S, h \rangle$:

$S^i = (Z \times \Pi \times N \times N - 1)\ \forall i$ so $s^i = (z^i, \pi^i, n^i, k^i)$.

$h : S \to Z$ is defined for the winner and the loosers as:

$$h_w = z^{k^w}_{\pi^N \circ .. \circ \pi^2 \circ \pi^1(w)} \qquad \text{(the winner)}$$

$$h_{i \neq w} = \begin{cases} z^{k^w}_{\pi^N \circ .. \circ \pi^2 \circ \pi^1(i)} \text{ if } z^{k^w} = z^w \\ 0 \text{ otherwise} \end{cases} \qquad \text{(the loosers)}$$

where $w = \sum_i n^i{}_{\mathrm{mod}\, N}$

The idea of eliminating envy via permutations is borrowed from "Divide and Permute" (Thomson (2005)). There however, only two players suggest divisions, while all others play permutations. In our games, N proposals instead of two may seem somewhat redundant. However, equal strategy spaces seem more appropriate for an (ex ante) equal treatment of players. As a more substantial innovation, we introduce a modulo game which awards the winner immunity to punishments. This is substantial as it yields better dynamic properties: it can easily be seen that the best reply dynamics (is likely to) produces

[4] If w (or k^w) could choose any reference allocation $z \in Z$, he would propose a division whereby only one share is different than zero, allocating it to himself with an opportune permutation. All players would then prefer to be winners (or k^w), so there would be no pure strategy equilibrium.

cycles in Divide and Permute, while in our mechanism it yields convergence. More details are in Section 4.

We have the following

Lemma 1 *Game form Γ^{EF} fully implements the EF correspondence in pure Nash equilibria, that is:*

i) $s \in pNE \Rightarrow h(s) \in EF$

ii) $z \in EF \Rightarrow \exists s \in pNE : h(s) = z.$

Proof. To prove $pNE \Rightarrow EF$, suppose $s \in pNE$, $h(s) = z$. It must then be the case that $z^{k^w} = z^w$. For if it were otherwise, any $j \neq w$ could conveniently deviate by playing a) $n^j : \sum_i n^i{}_{\bmod N} = w = j$; b) $k^j = j + 1$; c) a π^j that gives him a non-zero share (which exists in z^{j+1} because the z^is are partitions of ω). As a consequence, the range of shares that each player can reach by deviating contains $\rho = \{z_i^w\}_{i=1..N}$. Then, if $u_i(s)$ is the utility that i receives from share s, in equilibrium it must be $h_i = \arg\max_r u_i(z_r^w)$ $\forall i$. Hence, the outcome $h_i = z^w_{\pi^N \circ .. \circ \pi^2 \circ \pi^1(i)}$ is envy-free.

To prove $EF \Rightarrow pNE$ suppose $z \in EF$. The strategy profile $s^* : s^i = (z, id, 1, 1)$, with id the identity permutation, gives division z. Profile s^* is also a pNE: in fact: i) facing s^{*-i}, the range of shares attainable by i is $\{z_1 ... z_N\}$; ii) because $z \in EF$, profile s^* allocates to each player his preferred shared in this range. That is, s^* is a profile of mutual best replies. ■

3.2 Pareto optimality: Γ^{PO}

We now present a game that implements PO_+ allocations, that is Pareto optimal allocations where nobody receives a zero share. PO_+ allocations are clearly a superset of fair allocations, and some of them may generate envy. We don't worry about these latter, though, as they will be ruled out by Γ^{EF}, so we are not interested in their implementation. Again, we first describe the mechanism informally, giving it a flavour of a sequential game to illustrate its logic.

The key to optimality is that every Nash outcome will be a competitive equilibrium for some appropriate price vector. Welfare-theorem arguments then ensure Pareto efficiency. A little more in detail, each player i makes a proposal, i.e. suggests an allocation/price vector $d^i = (z^i, p^i)$. The mechanism selects one particular such d (we say in a moment *how* this is done). The selected proposal naturally defines N budget sets B_i, $i = 1..N$, one for each player. The game is built in such a way that each player i can obtain his favourite share within B_i. Thus, in equilibrium *all* i maximize within B_i, i.e. a competitive equilibrium results. Like in Γ^{EF}, two elements are essential in the construction: a) which proposal is selected? b) how can players choose their preferred shares? As for a), each player names an integer to be used in a modulo game. Also, each i indicates another player. The selected d is the one suggested by the player named by w, the winner of the modulo game. As for b), each player i calls a "reservation share" $\overline{z}^i$, which he does receive if both: i) he wins the modulo game and ii) $\overline{z}^i$ is "reasonable", that is if $\overline{z}^i \in B_i\left(d^{k^i}\right)$. If instead w claims for himself a share $\overline{z}_i \notin B_w\left(d^{k^w}\right)$, he's punished with a zero share. As for the loosers of the modulo game, they receive what is prescribed for them by d^{k^w}, but only if player k^w's proposal accomodates both i) w's proposal d^w and ii) claim $\overline{z}^w$; otherwise, they are punished with a zero share. The reason of the first punishment is clear: were the winner unrestrained, every i would try to be winner and there would be no equilibrium. The reason for the second punishment (on the loosers of the modulo game) will become apparent in the proof of Lemma 2.[5]

To formally describe the game we need more notation. Define D as the set of allocation-price proposals whereby every i gets a non-zero bundle:

$$D = \left\{(z,p) \in Z \times \mathbb{R}^l_{++} : p^T z_i > 0\right\}$$

For a given proposal $d = (z,p)$, the budget set of player i is:

$$B_i(d) = \left\{x \in \mathbb{R}^l_+ : 0 \le x \le \omega,\ p^T x \le p^T z_i\right\}$$

We can now describe the mechanism:

Game form Γ^{PO} is $\langle S, h\rangle$:

$S^i = D \times Z_i \times N \times (N-1)$ $\forall i$ so $s_i = \left(\left(p^i, z^i\right), \overline{z}_i, n^i, k^i\right)$.

[5] We might impose such punishment on k^w only. Lemma 2 and all subsequent results would still hold, but the definition of the outcome functions would take one extra line.

$h : S \to Z$ is defined for the winner and the loosers as:

$$h_w = \begin{cases} \overline{z}_w \text{ if } \overline{z}_w \in B_w\left(d^{k^w}\right) \\ 0 \text{ otherwise} \end{cases} \quad \text{(the winner)}$$

$$h_{i \neq w} = \begin{cases} z_i^{k^w} \text{ if } d^{k^w} = d^w \text{ and } \overline{z}_w = z_w^w \\ 0 \text{ otherwise} \end{cases} \quad \text{(the loosers)}$$

Lemma 2 *Game form Γ^{PO} fully implements PO_+ partitions in pNE, that is:*
i) *$s \in pNE \Rightarrow h(s) \in PO_+$*
ii) *$z \in PO_+ \Rightarrow \exists s \in pNE : h(s) = z$.*

Proof. To prove $pNE \Rightarrow PO_+$ we show that, in equilibrium, each i is getting his favourite share within his budget balance B_i (and these B_is are defined by a unique allocation and price vector). First observe that in equilibrium there is no $i : h_i = 0$ (any i can win the modulo game in order to pick a particular d^k and any share $\overline{z}_i \in B_i\left(d^k\right)$). Thus, $d^{k^w} = d^w$ and $h_i = z_i^w$ $\forall i$. As a consequence, the range of allocations over which any i is maximizing in equilibrium is $\rho_i = \{0\} \cup_{j \neq i} B_i\left(d^k\right)$. Because at least two players are making the same suggestion ($d^{k^w} = d^w$), we have $B_i\left(d^{k^w}\right) \subseteq \rho_i$ $\forall i$. Thus, a Nash outcome *a fortiori* maximises u_i within $B_i\left(d^{k^w}\right)$ $\forall i$. As a consequence, $h_i = z_i^{k^w}$ $\forall i$ is a competitive equilibrium with prices p^{k^w} and initial endowments $z_i^{k^w}$ $\forall i$ (and zero-trade).

To prove $z \in PO_+ \Rightarrow \exists s \in pNE : h(s) = z$, suppose $z \in PO_+$. By the second welfare theorem, there is a price vector p that supports z as a competitive equilibrium. The strategy profile $s^* : s^i = (z, p, 1)\, \forall i$ yields z and is a pNE. To see this latter fact, note that the range of allocations reachable for i is $\{0\} \cup B_i(z, p)$; because z^* is assumed to be a competitive equilibrium, every person optimizes within his budget, i.e. s^* is a profile of reciprocal best replies. ■

Note the role of the punishments on the loosers. We cannot allow the winner to name himself as proposer of d -otherwise any i would try to be winner so there would be no equilibrium. Thus, we assign to i a range of attainable shares $\rho_i = \{0\} \cup_{k \neq i} B_i\left(d^k\right)$. Then, if there were equilibria in which all proposals are different, there would *not* exist a single $d : B_i(d) \subseteq \rho_i$ $\forall i$. As a consequence, players would *not* be maximizing within a collection $\{B_i(d)\}_i$ defined by a *unique* price vector/endowments, as in a competitive equilibrium. Introducing a punishment when $d^{k^w} \neq d^w$, we force the.existence of at least two identical proposals, which in turn ensures that such a $\{B_i(d)\}_i$ does exist. On the other hand, punishing the case $\overline{z}_w \neq z_w^{k^w}$ is the simplest way to respect a physical constraint. Infact: in our construction d^{k^w} must become the "reference proposal"; now, if we are to give $\overline{z}_w$ to the winner we can distribute $z_i^{k^w}$ to all loosers i only if $\overline{z}_w = z_w^{k^w}$ (indeed: $\omega - \sum_{i \neq w} z_i^{k^w} = z_w^{k^w}$). Thus we must avoid $\overline{z}_w \neq z_w^{k^w}$; the simplest means to this is a punishment like the one above.

Finally: we could simplify the game eliminating the reservation proposal $\overline{z}_i$. Player w could be allowed to receive any $z_w \in B_w(.)$ by naming it in his d^w, instead of indicating it separately as "reservation share". The loosers would then be punished if $d^{k^w} \neq d^w$. A reservation proposal however makes it easier to prove the dynamic properties shown later on.

A few remarks put Γ^{PO} in relation with Thomson (2005). The idea of reaching efficiency via a competitive equilibrium is borrowed from there. However, like for Γ^{EF}, we have a game with equal strategy spaces to ensure an (ex-ante) equal treatment of players. Other radical changes are introduced to obtain better dynamic properties (it is easy to show that Thomson's efficiency game is likely to cycle for the best reply dynamics).

3.3 Fairness: Γ^F

The following mechanism combines elements of Γ^{EF} and Γ^{PO} to fully implements fair divisions. As in the previous games, a fair allocation emerges as a proposal made by (at least) two players. More in detail, each player: i) names a proposal d^i and a reservation share $\overline{z}_i$, ii) names an integer $n \in N$, , iii) indicates another proposal k^i, iv) chooses a permutation π^i. The winner of the modulo selects one proposal, to be taken as reference allocation (proposal d^{k^w}). Then, the outcome function allows any player to pick any bundle within $B_w\left(d^{k^w}\right)$ or, if he prefers, any share appearing in z^{k^w}.[6] As in Γ^{PO}, we force

[6] In equilibrium, w's best share in B_w will be already present in z^{k^w}, and is reached via permutations.

$d^{k^w} = d^w$ in equilibrium by punishing the loosers if this is not the case. As a result, players maximize within a family of budget balances defined by a unique price vector, so the reference allocation z^{k^w} is efficient. Permutations then ensure that each player can reach any share in z^{k^w}, thus ensuring envy-freeness. A formal description follows.

Game form Γ^F is $\langle S, h \rangle$:

$S^i = D \times Z_i \times N \times (N-1) \times \Pi\ \forall i$ so $s_i = ((p^i, z^i), \overline{z}_i, n^i, k^i, \pi^i)$.

$h : S \to Z$ is defined for the winner and the loosers as:

$$h_w = \begin{cases} \overline{z}_w \text{ if } \overline{z}_w \in B_w\left(d^{k^w}\right) \text{ and } d^{k^w} \neq d^w \\ z^{k^w}_{\pi^N \circ .. \circ \pi^2 \circ \pi^1 (w)} \text{ otherwise} \end{cases} \quad \text{(the winner)}$$

$$h_{i \neq w} = \begin{cases} z^w_{\pi^N \circ .. \circ \pi^2 \circ \pi^1 (i)} \text{ if } d^{k^w} = d^w \\ 0 \text{ otherwise} \end{cases} \quad \text{(the loosers)}$$

Recalling Lemma 1) and Lemma 2), the following is simply shown:

Proposition 1 *Game form* Γ^F *fully implements* F *partitions in* pNE.

Proof. To prove $pNE \Rightarrow F$, observe first that in a pNE it must be $d^{k^w} = d^w$ so that $h_i = z^w_{\pi^N \circ .. \circ \pi^2 \circ \pi^1 (i)} \forall i$. For if it were otherwise, k^w would be better off by playing $d^{k^w} = d^w$ (all other players too would rather deviate, to win the modulo game and choose opportune strategies). Also, because any player can be winner and receive $\overline{z}_w \in B_w\left(d^{k^w}\right)$, in equilibrium it must be the case that *every* player is maximizing within $\underset{j \neq i}{\cup} B_i\left(d^j\right) = \underset{j=1...N}{\cup} B_i\left(d^j\right) \supseteq B_i\left(d^{k^w}\right)$. Thus, h_i is a competitive equilibrium allocation with prices/endowments $\left(p^{k^w}, z^{k^w}\right)$, i.e. is efficient. Also, thanks to the permutations, $h_i = z^w_{\pi^N \circ .. \circ \pi^2 \circ \pi^1 (i)}$ maximizes u_i over $\{z^{k^w}_j\}_{j=1..N}$ for each i, i.e. it's envy-free.

As for the reverse inplication, consider a particular $z \in F$. Being $z \in PO_+$, there is a positive price vector p that sustain it as a competitive equilibrium. Then, it is immediate that strategies $s^i = (z, p, z_i, 1, 1, id)$ are a pNE producing outcome z. ■

We now turn to the simplicity properties by Saijo et al (1996) cited in the Introduction. It's immediately checked that Γ^F respects balancedness, best response and finite-dimension of strategy spaces. Forthrightedness requires that in equilibrium each player receives what he suggests. Thus, this property is satisfied by our game up to a permutation at worst (for any fair z, there is a forthrighted equilibrium where each i announces z plus the null permutation). Apart from Thomson (2005), the other "simple" game to implement fair divisions (see Introduction) is in Saijo et al.(1999). There, players only announce two quantities and prices. Our mechanism may seem less simple because it require players to suggest: i) full allocations, which is more than *two* bundles *plus* ii) additional elements like permutations and integers. However two objections should be raised: a) if the only requirements on the outcome function are feasibility, best response and forthrightedness, then our games are no more complex than Saijo et al.(1999)'s: m-dimensional strategies can be summarized into n-dimensional messages ($n < m$) by means of space-filling curves. b) if, on the other hand, we care about additional aspects of "simplicity", then our mechanism is arguably simpler than Saijo et al.(1999)'s. In particular, best replies can be immediately computed for our games, while they involve complex computations in Saijo et al.(1999). This is a crucial advantage when studying the dynamic properties of the mechanism.

4 Dynamic implementation

As argued in the Introduction, bare implementation is somewhat unsatisfactory for our problem: given the (necessary) multiplicity of equilibria of a mechanism, it is not clear if/how any of them can be coordinated upon. This is not a problem for game form Γ^F though, because its equilibria can be "learned". More precisely: a perturbed version of the best reply dynamics converges to the set of ε-equilibria with probability one. In turn, ε-equilibrium outcomes are "ε-fair" in a sense to be clarified soon. The necessity to consider ε-equilibria (and ε-fair outcomes) will become apparent in due course.

Given $\varepsilon > 0$ and strategy profile $\widehat{s}$ the set of ε-best replies of player i is:

$$BR_\varepsilon(\widehat{s}) = \{s^* \in S^i : u_i[h_i\left(s^*, \widehat{s}^{-i}\right)] \geq u_i[h_i\left(s^i, \widehat{s}^{-i}\right)] - \varepsilon \ \ \forall s^i \in S^i\}$$

Accordingly, we define $BR_0(.) = BR(.)$ the set of pure best replies.[7] With

[7] The fact we consider pure strategies only is not a limitation, as we show later on.

this notation, the set of ε-equilibria is

$$NE_\varepsilon = \{s : s^i \in BR_\varepsilon(s^{-i})\ \forall i\}$$

As anticipated, we will obtain convergence to ε-equilibria. What are then the fairness properties of the corresponding outcomes? We have the following

Lemma 3 *ε-equilibrium outcomes of Γ^F belong to the following set F_ε of "ε-fair" allocations:*

$$F_\varepsilon = \{z \in Z : [\nexists z' \in Z : u_i(z'_i) \geq u_i(z_i) + \varepsilon\ \forall i] \wedge [u_i(z_i) \geq u_i(z_j) - \varepsilon\ \forall i, j]\}$$

Observe that F_ε can rightly be said the set of "ε-fair" allocations because there: i) no gain is possible for everyone, that exceeds ε; ii) no individual is more than ε-envious. Also, Also, the set F_ε evidently shrinks to F as $\varepsilon \to 0$. Reminding the proof of Proposition 1, establishing the above Lemma is immediate:

Proof. If z is an ε-equilibrium outcome then $u_i(z_i) \geq u_i(z_j) - \varepsilon \forall i, j$ (evident because i can get j's share via permutations). On the other hand, suppose $\exists z' : u_i(z'_i) \geq u_i(z_i) + \varepsilon\ \forall i$. Because each i can obtain any share in $B_i(z, p)$, it is $pz'_i > pz_i\ \forall i$. Summing over i we get $p\omega = pz' > pz = p\omega$. ■

We introduce a last assumption; we comment it along the dynamics:

A1) the set of possible proposals $D = \{.., d, ..\}$ is finite and at least one of features a $z \in F_\varepsilon$.

4.1 Dynamics' specification

We call $s(t) = \left(s^1(t), s^2(t) .., s^N(t)\right)$ the strategy profile played at time t. To save notation, we use $BR(s^{-i}(t))$ or $BR^i(t)$ interchangeably, or even $BR(t)$ when there is no doubt about index i. Consider then the following discrete-time adjustment process:

D1) At each t, with probability $p^i \in (0, 1)$ player i updates his strategy;

D2) When i updates at t,

i) if $s^i(t-1) \in BR_\varepsilon(t-1)$ then $s^i(t) = s^i(t-1)$

ii) otherwise, any $s^i \in BR(t-1)$ is played with prob.> 0.

In words: at each time t, a random draw with probability p^i decides whether player i can update his strategy (the only constraint imposed on p^i is that it be uniformly bounded in time within $(0,1)$). When the random draw gives to i the possibility to update at t, player i chooses any best reply to the observed action profile $s^{-i}(t-1)$; however, he does not change his strategy if this turns out to be already an ε-best reply. Note that under **A1)** the set $BR^i(.)$ is finite (because finite is S^i), thus **D2** ii) is legitimate.

The set of rest points is

$$RP = \{s : prob\,(s\,(t+1) = s \mid s\,(t) = s) = 1\}$$

It should be clear that the set of rest points coincides with that of ε-Nash equilibria: $RP = NE_\varepsilon$. We will prove that RP is also an attractor.

We now comment on **D2)** and **A1)**. An equilibrium is reached when an equilibrium allocation (within an appropriate proposal d) is suggested. In turn, by implementation only fair allocations can appear in equilibrium. Now, because new proposals are put forward randomly (by **D2** ii), if players update strategies whenever they have an improving deviation (instead of an at least ε-improving one), an equilibrium would never emerge: the set F has a zero measure in Z so fair proposals are too "rare" to be hit by our players. We have two ways out: either we ensure some form of monotonicity (which amounts to new assumptions on the adustment process), or we give the target equilibrium set some mass. We choose the second way, and we do this considering F_ε, the set of ε-Nash outcomes. As we discussed above, it is appropriate to label F_ε the set of ε-Nash outcomes, because these latter are "ε-fair". The reason to introduce **A1)** is essentially the same motivating **D2)**. In equilibrium we need w and k^w to make the same proposal d. Thus, if ds were drawn from an infinite pool, no particular d could be agreed upon. The finiteness assumption could be disposed of if players had some preference for agreeing. To keep the assumptions on dynamics simple and general we prefer to introduce **A1)** -which seems quite reasonable anyway. Finally, note that **A1)** does not make **D2)** redundant. Infact, if we insisted on fairness (instead of ε-fairness) we would run into the problem that there exist no finite set of allocations that includes a fair one *whatever the preference profile,* while this is what is needed for implementation.

Finally, we put **D1)** and **D2)** in relation with the literature. **D1)** allows any number of players to update at t; a different assumption appears for example in Kim and Sobel (1995), where updating is sequential (assumption I). As it will be clear from the proof of Proposition 2, we only require that sequential updating be *possible*. As for **D2)**, it does *not* rules out the adoption of non-best replies. Similar formulations appear in Cabrales (1999), who allows for less-than-ε-improving updatings, but assumes that a non-improving strategy is never adopted (assumption Y3). It can be noted that there too it was necessary to consider ε-equilibria to obtain sure convergence.

4.2 Convergence to fair divisions

The focus of this work is on fair divisions. Although similar claims can be proven for game forms Γ^{EF} and Γ^{PO}, we only give the main result about Γ^{F}.

Proposition 2 *Under A1, D1 and D2, mechanism* Γ^F *is such that*
$\forall s(0)$, $prob\,[s(t) \in RP] \to 1$ *as* $t \to \infty$.

In words: whatever the starting point, sooner or later the process stops. Also, because $RP = NE_\varepsilon$ the limit allocation are in the set of "ε-fair" divisions.

The intuition behind convergence is simple. For any s^{-i} and any division z, there is a best reply whereby z is suggested (the winner's payoff is not affected by his own d). Thus, if the updating process goes on for a sufficiently long time, an ε-fair division (perhaps up to permutations) will eventually be suggested. At this point, provided that some conditions hold on current proposals, no player will gain from upsetting z as reference division. Instead, it will be sufficient that permutations are adjusted until the shares are allocated without envy. The resulting strategies will be an ε-equilibrium, so a rest point. The following proof puts this in formal terms.

Proof. of Proposition 2

The winner's payoff does not depend on his own d. On the other hand, the winner's range of attainable shares is the largest. Thus, it is always a best reply to win the modulo game and make a *new* proposal d. As a consequence, any profile of proposals $\{d^i\}^{i=1..N}$ in D^N will emerge before some time t, with probability tending to 1 as $t \to \infty$. In particular, a profile d^* will be put forward containing a z^* such that

i) $\forall i, j$ it is $u_i(z_i^*) + \varepsilon \geq u_i(z') \;\; \forall z' \in B_i(d^j)$;

ii) if $u_i(z_j^r) + \varepsilon \geq u_i(z_i^*)$ for some i, j, r then there is no $k^i = r$.

By assumption **A1**, proposals under i)-ii) do exist in D^N: any $\{(z,p), (z,p), .., (z,p)\}$, with $z \in F_\varepsilon$ and p an appropriate supporting price vector satisfies i)-iii) independently of the played strategies.[8]

Condition i) ensures that, if d^* has emerged, there is no advantage for the winner to select a different proposal d^j -by so doing, he could reach some $z' \in B_i(d^j)$. On the other hand, ii) says that if division $z^r \neq z^*$ is more attractive for i than z^* (because it contains a certain share z_j^r), then z^r cannot become a reference division, because no one named player r (if player k had named player r, i could make k win to have z^r adopted as reference allocation).

Once d^* is put forward, it is sufficient that the next updater chooses $d^i = d^{i^*}$ and $\pi^i : \pi^N \circ ...\pi^i \circ ... \circ \pi^1(.) = id$, and the outcome is z^*. Conditions i)-ii) ensure that the resulting strategy profile is a ε-Nash equilibrium (infact, no player, has a convenient deviation, because z_i^* is the best share for each i, in the range of attainable shares). As a consequence the given strategy profile is a rest point. ■

Remark 1 (Mixed Strategies) *Our games yield full implementation in pure-strategy equilibria. However, implementation fails when mixed strategy equilibria are used. This is a fact known to happen in many classic implementation mechanisms. Consider for example Γ^{EF} and the case of two players. The strategy profile where each player: i) assigns everything to one share, ii) randomly choose between the swap and the id permutation with probability $\frac{1}{2}$, is clearly a NE. However, its outcome is non-EF with probability $\frac{1}{2}$. If we look at the dynamics though, we still find that the limit outcomes are fair. The reason is simple. Proposition 2) shows that a rest point does emerge; then, because players*

[8] There might exist others such ds, too. We don't go into details, but it is clear that the finer is the grid of possible allocation/prices, the more likely it is that there will be many such proposals, making convergence faster.

best respond to observed strategies (not to mixtures), if $\sigma = (\sigma^1, ..., \sigma^N) \in RP$ *then each* σ^i *must be a best reply against each action profile in the support of* σ *itself. It follows that any realized outcome from a rest point is an equilibrium outcome, which is fair by implementation.*

5 Conclusions

The focus of this paper is on a classic fair division problem: allocating homogeneous divisible goods among players, in a way that is both envy-free and efficient. While the existence of these divisions is a well established fact, the implementation side of the problem has been less thoroughly studied.

This paper gives a procedure with the following properties: a) the procedure is a formally defined game form; we look at its Nash equilibria; b) strategy spaces are symmetric and relatively simple: essentially players are requested to announce division proposals and price vectors. Quite straightforward is also the outcome function: best replies are immediately calculated for any strategy profile. Most importantly c) the equilibria are learned by players who update their strategies according to best-reply dynamics. Features a) and b) partly appear in Thomson (2005), to which this paper is inspired. Feature c) is new and, I suggest, important for a mechanism to be meaningful solution to an applied implementation problem, especially in the presence of multiple equilibria.

The mechanisms presented here are tailored to a specific kind of dynamics. Although some rationality is embodied in the concept of Nash equilibrium, the considered dynamics are definitely myopic. Thus, one might doubt that players choose at t best replies to $t-1$ observed actions. For example, if the shares are given out only when a rest point is reached, why should we exclude that a Nash equilibrium is deviated from? Even if a deviation from a NE is not immediately convenient, it might trigger further deviations by other players, so the final result might be a net gain for the first deviator. Similarly, one might be interested in looking at coalitional equilibrium concepts. These considerations evidently lead away from the Nash equilibrium concept.

In keeping with the classic implementation theory, this paper concentrates on Nash equilibria, and in particular on pure-strategy equilibria. However, mixed strategy equilibria as well are taken care of. In particular, mixed strategies with (random) non-fair outcomes cannot be rest point; this result descends directly from the assumptions on dynamics. If we do not accept such specification of dynamics, mixed strategy could become relevant again either when i) players observe mixed strategies -which would not be very realistic either, or when ii) players form complex beliefs about their opponents' strategies. In the latter case, we should again abandon Nash implementation, and look at appropriately defined sequential equilibria.

References

[1] Brams S.J., Taylor A.D. (1995), "An Envy-Free Cake Division Protocol", *American Mathematical Monthly*, 102 (1), 1995.

[2] Brams S.J., Taylor A.D. (1996), "Fair Division: From Cake-Cutting to Dispute Resolution", CUP 1996.

[3] Cabrales A. (1999), "Adaptive Dynamics and the Implementation Problem with Complete Information", *Journal of Economic Theory*, 86, 1999.

[4] Cabrales A., Ponti G. (2000), "Implementation, Elimination of Weakly Dominated Strategies and Evolutionary Dynamics", *Review of Economic Dynamics*, 3, 2000.

[5] Dutta B; Sen A., Vohra R. (1995), "Nash Implementation through Elementary Mechanisms in Economic Environments", *Economic Design*, 1(2) 1995.

[6] Kim Y. G., Sobel J., (1995), "An Evolutionary Approach to Pre-play Communication", *Econometrica*, 63(5) 1995.

[7] Maskin E. (1977), Nash Equilibrium and Welfare Optimality, MIT mimeo.

[8] Maskin E., Sjostrom T. (2002), "Implementation Theory", *Handbook of Social Sciences and Welfare*, vol.1, North Holland.

[9] Postlewaite A., Wettstein D. (1989), "Feasible and Continuous Implementation", *The Review of Economic Studies*, 56 (4), 1989.

[10] Saijo T., Tatamitami Y. and Yamato Y. (1996), "Toward Natural Implementation", *International Economic Review,* 37 (4), 1996.

[11] Saijo T., Tatamitami Y. (1999), "Characterizing Natural Implementability: The Fair and Walrasian Correspondences", *Games and Economic Behavior* 28, 1999.

[12] Sandholm W.H. (2002), "Evolutionary Implementation and Congestion Pricing", *Review of Economic Studies,* 69, 2002.

[13] Sjöström, T (1994), "Implementation in Undominated Nash Equilibria without Integer Games", *Games and Economic Behavior* 6(3) 1994.

[14] Thomson W. (2005), *Games and Economic Behavior*, Vol. 52, Issue 1, July 2005, Pages 186-200.

[15] Varian, H. R. (1974) "Equity, Envy, and Efficiency", *Journal of Economic Theory* 9, 63-91.